Vicinity-Driven Paragraph and Sentence Alignment for Comparable Corpora

Gustavo Henrique Paetzold and **Lucia Specia**
Department of Computer Science
University of Sheffield, UK
{g.h.paetzold,l.specia}@sheffield.ac.uk

arXiv:1612.04113v1 [cs.CL] 13 Dec 2016

Abstract

Parallel corpora have driven great progress in the field of Text Simplification. However, most sentence alignment algorithms either offer a limited range of alignment types supported, or simply ignore valuable clues present in comparable documents. We address this problem by introducing a new set of flexible vicinity-driven paragraph and sentence alignment algorithms that 1-N, N-1, N-N and long distance null alignments without the need for hard-to-replicate supervised models.

1 Introduction

Early Text Simplification (TS) approaches relied heavily on manually-crafted simplification rules (Chandrasekar et al., 1996; Siddharthan, 2014) due to a lack of resources from which to automatically learn or extract them. Complex-to-simple parallel corpora effectively changed the nature of contributions pertaining to TS upon their arrival.

The most widely used resource of this kind is the Wikipedia and Simple Wikipedia parallel corpus (Kauchak and Barzilay, 2006). Using this corpus, authors have devised Syntactic Simplification approaches that learn simplification rules from aligned sentences (Siddharthan, 2006; Woodsend and Lapata, 2011), Machine Translation approaches that translate from complex to simple text (Zhu et al., 2010; Bach et al., 2011), Lexical Simplifiers that extract complex-to-simple word correspondences from word alignments (Biran et al., 2011; Horn et al., 2014), and even hybrid lexico-syntactic simplifiers (Paetzold and Specia, 2013; Feblowitz and Kauchak, 2013). Other examples of parallel corpora that have been used in TS are the medical-domain corpus of (Deléger and Zweigenbaum, 2009), and the newly introduced Newsela corpus (Xu et al., 2015).

Most of these corpora come, however, as a set of complex-to-simple articles aligned at document level. In order for them to be used by Machine Translation systems, Lexical Simplification strategies and Tree Transduction frameworks, for example, these documents need to be aligned at sentence level.

Various strategies have been created to address this problem. Barzilay and Elhadad (2003) present an approach that first aligns paragraphs with sophisticated clusterers and classifiers learned from manually annotated instances, then produces sentence alignments through a dynamic programming algorithm that finds an alignment path between the sentences in aligned paragraphs. Coster and Kauchak (2011) use the same sentence alignment approach, but opt instead for a much simpler paragraph alignment algorithm: they simply align any pair of paragraphs of which the TF-IDF similarity is higher than a manually set threshold of 0.5. Similarly, Bott and Saggion (2011) also make use of an elaborate, difficult to replicate supervised approach in order to learn the likelihood of two portions of text being aligned. Smith et al. (2010) bypass the step of paragraph alignment and use instead a supervised Conditional Random Fields ranker for sentence alignment directly.

Although these algorithms have been proven useful, they do not exploit the comparability between the aligned documents' paragraphs and sentences, hence the need for elaborate supervised classifiers and rankers. The sentence aligner used by both Barzilay and Elhadad (2003) and Coster and Kauchak (2011) suffers from yet another limitation: it does not allow for multiple consecutive sentence skips, which can lead to incorrect alignments in scenarios where the number

of sentences in a pair of aligned paragraphs is not similar.

In this paper, we introduce a new set of paragraph and sentence alignment algorithms designed to address these limitations.

2 Clues in Comparable Corpora

By inspecting a handful of documents from the Wikipedia-Simple Wikipedia and the Newsela corpus, we noticed that they are divided in numerous paragraphs, and that the order with which the information is presented is consistent throughout different versions of the same article. This means that, most of the time, one can safely assume that, if the ith and jth paragraph of a pair of documents is aligned, then the next alignment (x, y) will be such that $x \geq i$ and $y \geq j$.

Since paragraphs offer valuable alignment clues, we believe that algorithms that bypass paragraph alignment, such as that of Smith et al. (2010), are not suitable for this task. We also noticed that the order of the information within aligned paragraphs is also consistent, but that often times aligned paragraphs will have a noticeably large disparity in number of sentences they have. Consequently, sentence alignment algorithms that do not allow for multiple consecutive skips, such as that of Barzilay and Elhadad (2003), may not be suitable for this task either.

Xu et al. (2015) produced sentence alignments for the Newsela corpus by simply calculating the Jaccard similarity score of each pair of sentences in a set of documents, then selecting any pairs that achieve a similarity score above a manually set threshold. This approach, however, suffers from several limitations also, since it does not exploit the aforementioned observations in any capacity, nor does it allow for N-1 or N-N alignments.

We address this problem by introducing new, flexible paragraph and sentence alignment algorithms that better document-aligned corpora. We describe them in what follows.

3 Paragraph Alignment Algorithm

In order to produce paragraph alignments, we employ a search algorithm that allows 1-1, 1-N, N-1, N-N and null alignments (i.e. unaligned paragraphs between alignments).

Algorithm 1 exploits a matrix $M \in \mathbb{R}^{\|D_1\| \times \|D_2\|}$ where $\|D_1\|$ represents the number of paragraphs in a document D_1, $\|D_2\|$ the number of paragraphs in a document D_2, and $M(i, j)$ the similarity between the ith paragraph in D_1 and the jth paragraph in D_2. As a similarity metric, we use the maximum TF-IDF cosine similarity between all possible pairs of sentences in $[i, j]$ i.e. the similarity between two paragraphs is equal to that of the most similar pair of sentences in them. We choose this metric due to the fact that, even though the sentences in equivalent paragraphs of documents with different reading levels are usually very distinct in both form and vocabulary, often times paragraphs will share at least one very similar (or even identical) sentence between them, which is a strong indicator of a good paragraph alignment. These similar sentences are often so because of long quotes from subjects interviewed for the news articles.

The goal of Algorithm 1 is to find a path $A \in \mathbb{Z}^{\|D_1\| \times \|D_2\|}$ where $A(i, j) = 1$ if there is an alignment between paragraphs $[i, j] \in D_1 \times D_2$, and $A(i, j) = 0$ otherwise. The search for path A starts from the assumption that $A(1, 1) = 1$, or in other words, the first paragraph of D_1 and D_2 are aligned. We exploit this assumption because the first paragraph in most document-aligned corpora available refer to the articles' titles. The algorithm then initializes a control set of coordinates $[c_x, c_y] = [1, 1]$ that represents the point from which to search for the next alignment in A. The next alignment is then searched for in the first vicinity V1, which represents 1-1, 1-N and N-1 alignments. If there is not a paragraph pair $[n_x, n_y]$ in V1 such that $M(n_x, n_y) \geq \alpha$ then the next alignment is searched for in the second vicinity V2, which represents single paragraph skips. If no pair $[n_x, n_y]$ in V2 has $M(n_x, n_y) \geq \alpha$, then the algorithm searches for the the next pair $[n_x, n_y]$ with $M(n_x, n_y) \geq \alpha$ with the shortest euclidean distance to $[c_x, c_y]$ in the third and final vicinity V3, which represents long-distance paragraph skips. Finally, the update $[c_x, c_y] = [n_x, n_y]$ is made, and the alignment $A(c_x, c_y) = 1$ is added. This process is repeated until there are no more paragraphs left to be aligned, or until the algorithm finds no suitable $[n_x, n_y]$ to follow current alignment $[c_x, c_y]$ in vicinity V3. Notice that, although we use only three vicinities, the algorithm can be easily adapted to support as many distinct vicinities as suitable. After A is found, the paragraphs in all 1-N, N-1 and N-N alignments are concatenated

so that our sentence alignment algorithm can more easily search for equivalent sentences in them.

Algorithm 1: Paragraph Alignment

input: M, D_1, D_2, α;
output: A;

$A \leftarrow zeros\left(\|D_1\|, \|D_2\|\right)$;
$[c_x, c_y] \leftarrow [1, 1]$;
while $[c_x, c_y] \neq nil$ **do**
  $A(c_x, c_y) \leftarrow 1$;
  V1 $\leftarrow \{[c_x, c_y+1], [c_x+1, c_y], [c_x+1, c_y+1]\}$;
  V2 $\leftarrow \{[c_x+2, c_y+1], [c_x+1, c_y+2]\}$;
  V3 $\leftarrow \{[x, y] \mid l_x \geq x \geq c_x$ and $l_y \geq y \geq c_y\}$;
  $[n_x, n_y] \leftarrow [x, y]$ with best $M(x, y)$ in V1;
  if $M(n_x, n_y) < \alpha$ **then**
    $[n_x, n_y] \leftarrow [x, y]$ with best $M(x, y)$ in V2;
    if $M(n_x, n_y) < \alpha$ **then**
      $[n_x, n_y] \leftarrow [x, y]$ with $M(x, y) \geq \alpha$ in V3 with shortest euclidean distance to $[c_x, c_y]$;
    end
  end
  $[c_x, c_y] \leftarrow [n_x, n_y]$;
end

4 Sentence Alignment Algorithm

Our sentence alignment algorithm employs the same principles behind our paragraph alignment algorithm: it also allows for 1-1, 1-N, N-1 and null alignments, and works under the assumption that N-N alignments can be inferred from consecutive 1-1 alignments.

Algorithm 2 exploits a matrix $M \in \mathbb{R}^{\|P_1\| \times \|P_2\|}$ where $\|P_1\|$ represents the number of sentences in a paragraph P_1, $\|P_2\|$ the number of sentences in a paragraph P_2, and $M(i, j)$ the similarity between the ith sentence in P_1 and the jth sentence in P_2. As a metric, we use TF-IDF cosine similarity.

In search of path $A \in \mathbb{Z}^{\|P_1\| \times \|P_2\|}$, our algorithm first finds the initial alignment point $[c_x, c_y]$ with $M(c_x, c_y) \geq \alpha$ with the shortest euclidean distance to $[0, 0]$. Notice that this algorithm cannot exploit the assumption that the first sentences in a pair of paragraphs will be aligned. The next alignment $[n_x, n_y]$ is then searched for in the immediate vicinity V1. If $M(n_x, n_y)$ is smaller than a minimum α, then it searches for the the next pair $[n_x, n_y]$ with $M(n_x, n_y) \geq \alpha$ with the shortest euclidean distance to $[c_x, c_y]$ in vicinity V2. Otherwise, if $[n_x, n_y] = [c_x+1, c_y+1]$, then the new 1-1 alignment is added to A. However, if a 1-N ($[n_x, n_y] = [c_x, c_y+1]$) or N-1 ($[n_x, n_y] = [c_x+1, c_y]$) alignment is found, then the algorithm performs a secondary loop in order to find the size

Algorithm 2: Sentence Alignment

input: M, P_1, P_2, α, β;
output: A;

$A \leftarrow zeros\left(\|P_1\|, \|P_2\|\right)$;
$[c_x, c_y] \leftarrow [x, y]$ with $M(x, y) \geq \alpha$ with shortest euclidean distance to $[0, 0]$;
if $[c_x, c_y] \neq nil$ **then**
  $A(c_x, c_y) \leftarrow 1$;
end
while $[c_x, c_y] \neq nil$ **do**
  V1 $\leftarrow \{[c_x, c_y+1], [c_x+1, c_y], [c_x+1, c_y+1]\}$;
  V2 $\leftarrow \{[x, y] \mid l_x \geq x \geq c_x$ and $l_y \geq y \geq c_y\}$;
  $[n_x, n_y] \leftarrow [x, y]$ with best $M(x, y)$ in V1;
  case $M(n_x, n_y) < \alpha$ **do**
    $[n_x, n_y] \leftarrow [x, y]$ with $M(x, y) \geq \alpha$ in V2 with shortest euclidean distance to $[c_x, c_y]$;
    $[c_x, c_y] \leftarrow [n_x, n_y]$;
    $A(c_x, c_y) \leftarrow 1$;
  case $[n_x, n_y] = [c_x+1, c_y+1]$ **do**
    $[c_x, c_y] \leftarrow [n_x, n_y]$;
    $A(c_x, c_y) \leftarrow 1$;
  case $[n_x, n_y] = [c_x, c_y+1]$ **do**
    $A(n_x, n_y) \leftarrow 1$;
    $N \leftarrow 1$;
    while $M(n_x, n_y : n_y+N) > M(n_x, n_y : n_y+N-1) - \beta$ and $M(n_x, n_y : n_y+N) > M(n_x+1, n_y : n_y+N)$ **do**
      $A(n_x, n_y+N) \leftarrow 1$;
      $N \leftarrow N+1$;
    end
    $[c_x, c_y] \leftarrow [n_x, n_y+N-1]$;
  case $[n_x, n_y] = [c_x+1, c_y]$ **do**
    $A(n_x, n_y) \leftarrow 1$;
    $N \leftarrow 1$;
    while $M(n_x : n_x+N, n_y) > M(n_x : n_x+N-1, n_y) - \beta$ and $M(n_x : n_x+N, n_y) > M(n_x : n_x+N, n_y+1)$ **do**
      $A(n_x+N, n_y) \leftarrow 1$;
      $N \leftarrow N+1$;
    end
    $[c_x, c_y] \leftarrow [n_x+N-1, n_y]$;
  end
end

of N. N is incremented until the accumulated similarity for the concatenated sentences given the current size of N is smaller than the similarity for $N-1$ minus a slack value β, or until the similarity of the adjacent diagonal ($M(n_x+1, n_y : n_y+N)$ or $M(n_x : n_x+N, n_y+1)$) is larger than the accumulated similarity. This process is repeated until there are no more paragraphs left to be aligned, or until the algorithm finds no suitable new alignment candidates.

5 Final Remarks

We presented a new set of vicinity-driven paragraph and sentence alignment algorithms for

document-aligned corpora. Our algorithms addresses the limitations of previous approaches by exploiting clues in comparable documents that are often neglected. Unlike many early strategies, our algorithms allow for 1-N, N-1, N-N and long-distance null alignments.

In the future, we aim to conduct both intrinsic and extrinsic performance comparisons between our algorithms and earlier approaches. We also aim to employ our algorithms in the creation of new paragraph and sentence-aligned corpora to be made available to the public.

References

[Bach et al.2011] Nguyen Bach, Qin Gao, Stephan Vogel, and Alex Waibel. 2011. Tris: A statistical sentence simplifier with log-linear models and margin-based discriminative training. In *IJCNLP*.

[Barzilay and Elhadad2003] Regina Barzilay and Noemie Elhadad. 2003. Sentence alignment for monolingual comparable corpora. In *Proceedings of the 2003 EMNLP*, pages 25–32.

[Biran et al.2011] Or Biran, Samuel Brody, and Noémie Elhadad. 2011. Putting it simply: a context-aware approach to lexical simplification. In *Proceedings of the 49th ACL*, pages 496–501.

[Bott and Saggion2011] Stefan Bott and Horacio Saggion. 2011. An unsupervised alignment algorithm for text simplification corpus construction. In *Proceedings of the 2011 MTTG*, pages 20–26.

[Chandrasekar et al.1996] Raman Chandrasekar, Christine Doran, and Bangalore Srinivas. 1996. Motivations and methods for text simplification. In *Proceedings of the 16th COLING*, pages 1041–1044.

[Coster and Kauchak2011] William Coster and David Kauchak. 2011. Simple english wikipedia: A new text simplification task. In *Proceedings of the 49th ACL*, pages 665–669.

[Deléger and Zweigenbaum2009] Louise Deléger and Pierre Zweigenbaum. 2009. Extracting lay paraphrases of specialized expressions from monolingual comparable medical corpora. In *Proceedings of the 2nd Workshop on Building and Using Comparable Corpora: from Parallel to Non-parallel Corpora*, pages 2–10.

[Feblowitz and Kauchak2013] Dan Feblowitz and David Kauchak. 2013. Sentence simplification as tree transduction. In *Proceedings of the 2nd Workshop on Predicting and Improving Text Readability for Target Reader Populations*, pages 1–10.

[Horn et al.2014] Colby Horn, Cathryn Manduca, and David Kauchak. 2014. Learning a lexical simplifier using wikipedia. In *Proceedings of the 52nd ACL*, pages 458–463.

[Kauchak and Barzilay2006] David Kauchak and Regina Barzilay. 2006. Paraphrasing for automatic evaluation. In *Proceedings of the 2006 NAACL*, pages 455–462.

[Paetzold and Specia2013] Gustavo H. Paetzold and Lucia Specia. 2013. Text simplification as tree transduction. In *Proceedings of the 9th STIL*, pages 116–125.

[Siddharthan2006] Advaith Siddharthan. 2006. Syntactic simplification and text cohesion. *Research on Language and Computation*, 4(1):77–109, March.

[Siddharthan2014] Advaith Siddharthan. 2014. A survey of research on text simplification. *International Journal of Applied Linguistics*, pages 259–298.

[Smith et al.2010] Jason R Smith, Chris Quirk, and Kristina Toutanova. 2010. Extracting parallel sentences from comparable corpora using document level alignment. In *Proceedings of the 2010 HLT*, pages 403–411.

[Woodsend and Lapata2011] Kristian Woodsend and Mirella Lapata. 2011. Learning to simplify sentences with quasi-synchronous grammar and integer programming. In *Proceedings of the 2011 EMNLP*, pages 409–420.

[Xu et al.2015] Wei Xu, Chris Callison-Burch, and Courtney Napoles. 2015. Problems in current text simplification research: New data can help. *Transactions of the Association for Computational Linguistics*, 3:283–297.

[Zhu et al.2010] Zhemin Zhu, Delphine Bernhard, and Iryna Gurevych. 2010. A monolingual tree-based translation model for sentence simplification. *Computational Linguistics*, (August):1353–1361.